

METHODS AND SYSTEMS EMPLOYING DIGITAL WATERMARKINGField of the Invention

5 The present invention relates to applications of digital watermarking in conjunction with audio, video, imagery, and other media content.

Background

10 Watermarking (or "digital watermarking") is a quickly growing field of endeavor, with several different approaches. The present assignee's work is reflected in U.S. Patents 5,841,978, 5,768,426, 5,748,783, 5,748,763, 5,745,604, 5,710,834, 5,636,292, 5,721,788, and laid-open PCT applications WO97/43736 and WO99/10837. Other work is illustrated by U.S. Patents 5,734,752, 5,646,997, 5,659,726, 5,664,018, 5,671,277, 5,687,191, 5,687,236, 5,689,587, 5,568,570, 5,572,247, 5,574,962, 5,579,124, 5,581,500, 5,613,004, 5,629,770, 5,461,426, 5,743,631, 5,488,664, 5,530,759, 5,539,735, 4,943,973, 5,337,361, 5,404,160, 5,404,377, 15 5,315,098, 5,319,735, 5,337,362, 4,972,471, 5,161,210, 5,243,423, 5,091,966, 5,113,437, 4,939,515, 5,374,976, 4,855,827, 4,876,617, 4,939,515, 4,963,998, 4,969,041, and published foreign applications WO 98/02864, EP 822,550, WO 97/39410, WO 96/36163, GB 2,196,167, EP 777,197, EP 736,860, EP 705,025, EP 766,468, EP 782,322, WO 95/20291, WO 96/26494, WO 96/36935, WO 96/42151, WO 97/22206, WO 97/26733.

20 Most of the work in watermarking, however, is not in the patent literature but rather in published research. In addition to the patentees of the foregoing patents, some of the other workers in this field (whose watermark-related writings can be found by an author search in the INSPEC database) include I. Pitas, Eckhard Koch, Jian Zhao, Norishige Morimoto, Laurence Boney, Kineo Matsui, A.Z. Tirkel, Fred Mintzer, B. Macq, Ahmed H. Tewfik, Frederic Jordan, 25 Naohisa Komatsu, and Lawrence O'Gorman.

The artisan is assumed to be familiar with the foregoing prior art.

In the present disclosure it should be understood that references to watermarking encompass not only the assignee's watermarking technology, but can likewise be practiced with any other watermarking technology, such as those indicated above.

30 Watermarking has various uses, but the present specification details several new uses that provide functionality and features not previously available.

Brief Description of the Drawings

Fig. 1 is a diagram showing the participants, and channels, involved in the distribution of music.

Fig. 2 shows a conceptual model of how music artists, record labels, and E-Music distributors can all interact with a Media Asset Management System, of which several are detailed in the following specification.

Detailed Description

For expository convenience, much of the following discussion focuses on music, but the same principles and techniques are largely or wholly applicable to other source data, whether non-music audio, video, still imagery, printed materials, etc.

Music Asset Management

Referring to the figures, the music distribution process begins with a creative artist 10. The artist's music has traditionally been distributed by a record label 12. (While the following discussion refers to distribution through such a label, it should be understood that such distribution can just as well be effected directed under the artist's control, without a record label intermediary.)

In traditional distribution 14, the record label produces tangible media, such as records, tapes, videos (e.g. music videos), and CDs 16. These media are physically distributed to end-consumers 18. Additionally, the label 12 distributes the music media to outlets 20, such as radio and TV stations, cable and satellite systems, etc., which broadcast (or narrowcast) the artist's work to an audience. Distribution through such media outlets may be monitored by playout tracking services. Playout tracking data, collected by firms including Arbitron, Nielsen, ASCAP, BMI, etc., can be used to compute royalty payments, to verify broadcast (e.g. for advertising), etc.

Increasingly, the distribution of the music to the media outlets is performed electronically. Such distribution first took the form of analog audio over high quality landlines or satellite channels. Digital audio quickly supplanted analog audio in such distribution channels due to higher fidelity.

More recently, distribution of the music from the record labels to the media outlets has occurred over secure links, now including the internet. Such security was first provided simply by scrambling the audio signal or data. More sophisticated "container"-based systems are now coming into vogue, in which the audio is "packaged" (often in encrypted form) with ancillary data.

Electronic distribution of music to the consumer is also gaining popularity, presently in the MP3 format primarily. The music providers may deal directly with the public, but more commonly effect such consumer distribution through a newly emerging tier of digital media outlets, such as internet sites that specialize in music. From such sites, consumers can download digital audio files into personal digital audio players. (The Diamond Rio, and the Audible MobilePlayer devices are some of the first of what will doubtless be a large number of entrants into this personal internet audio appliance market.) Or the downloaded data can be stored by the consumer-recipient onto any other writeable media (e.g. hard disk, CD, DVD, tape, videotape, etc.). Typically a personal computer is used for such downloading, but this intermediary may be dispensed with by coupling next generation of personal audio appliances to an internet-like link.

The data downloaded by the consumer can be stored either in the native digital format, translated into another digital format (which translation may include decryption), converted into analog and recorded in analog form, etc.

Unauthorized copying or use of the music can occur anywhere in the foregoing channels. However, one of the greatest risks occurs once the music has been delivered to the consumer (whether by tangible media, by traditional broadcast media outlets, by emerging digital distribution, or otherwise).

The general idea of embedding auxiliary data into music (i.e. watermarking) has been widely proposed, but so far has been of limited applicability.

For example, GoodNoise is planning to embed a digital signature -- termed a multimedia identifier, or MMI -- in its MP3 music. MMI will register the song and its author with a licensing number. In addition to providing information about the songwriter and distributor, this digital encoding may also include lyrics, liner notes, and other information. But all of the proposed uses serve only to convey information from the distributor to the consumer; use for "tracking" is actively disclaimed. (Wired News, "GoodNoise Tags MP3 Files," February 3, 1999.)

The Genuine Music Coalition – a partnership of various companies in the music distribution business – likewise has announced plans to employ watermarking of MP3 music. The watermarking technology, to be provided by Liquid Audio, will convey data specifying the artist or producer contact, copyright data, and a number to track ownership. The Coalition hopes that the provision of this embedded information will help thwart piracy. Industry observers believe Liquid Audio will next introduce playback technology only plays audio in which its watermark is detected. (Wired News, “Liquefying MP3,” January 23, 1999.)

A similar initiative has been announced by the Recording Industry Association of America (RIAA). Termed the Secure Digital Music Initiative (SDMI), the program seeks to define a voluntary specification that will assure proper compensation to those who produce and distribute music. One element of the system will likely be a watermarking component. (Dow Jones Newswire, “Spurred By Maverick Technology, Music Industry Eyes Web,” December 31, 1998.)

Yet another initiative has been announced by Solana and ASCAP. Other companies promoting watermarking for music include Aris Technology, MCY.com, and AudioSoft.

The watermark payload can represent various types of data. An exemplary payload includes data relating to the artist, distribution entity, title, and copyright date/proprietor. Additionally, the payload can include a digital object identifier – an ISBN-like number issued by a central organization (e.g. a rights management organization) to uniquely identify the work.

Such payload data can be encoded literally (e.g. the title by a series of ASCII characters, etc.). In other embodiments, codes or abbreviations can be employed – with each code having a known meaning. In still other embodiments, the data can be meaningless by itself, but may serve as a key (e.g., a Unique Identifier, or UID) into a remote data database or repository. An example of such a remote data repository is a web site at a Master Global Address (MGA) associated with content, as detailed below.

An exemplary data payload may, for example, have the following format:

A	B	C	D	E	F	G	H	I		
---	---	---	---	---	---	---	---	---	--	--

Where A is a six-byte (8-bits to a byte) ASCII string serving as a digital object identifier (which may serve as a link to a Master Global Address through a default name server, as discussed

5

10

15

20

25

other fields, the streaming data field may not endlessly repeat the same data, but can convey data that changes with time.

Desirably, the encoding is performed in a manner permitting recovery of the watermark data even if the audio is corrupted, e.g. by format conversion, re-sampling, tape wow and flutter, compression, coding, or various forms of audio processing (e.g. filtering, pre-emphasis, re-scaling, etc.). One way to provide for such robustness is to encode a signal of known character that can be recognized through all such corruption. By identifying such known signal, the watermark signal can then be decoded. (The known signal can take various forms, e.g. a synchronization signal, a marker signal, calibration signal, a universal code signal as described in applicant's patents, etc.)

In some embodiments, a watermark "dial-tone" signal is provided. This dial-tone signal is a low amplitude, relatively wideband, repetitive signal that commonly conveys only limited information (e.g. a single bit of information). Its presence in an audio signal can serve as a "do not record," or similar instruction signal. Alternatively, or in addition, the dial-tone signal can serve as an aid in "locking" to a plural-bit digital watermark signal that is also encoded in the audio. For example, the cyclical repetition of the signal can serve to identify the start of the plural-bit digital watermark signal. Or the spectrum or repetition rate of the signal can identify any temporal corruption of the audio. An exemplary such signal is detailed as a "simple universal code" in Patent 5,636,292.

A track of music can be pre-authorized for specified types of use. For example, the usage control string of the watermark payload may include a six-bit field detailing the classes of devices for which the audio is authorized. Each bit would correspond to a different class of device. Class 1 devices may be personal playback devices with only analog-audio output. Class 2 devices may be personal entertainment devices capable of outputting music in digital (e.g. MP3, redbook, *.WAV) format, as well as analog audio. Class 3 devices may be personal computer systems (i.e. with essentially unlimited ability for processing and outputting digital audio). Etc., etc. A device to which such MP3 audio is provided would check the usage control string data to determine whether it is authorized to utilize the audio. A personal playback device with analog-only output, for example, would examine the first bit of the usage control string. If it was "1," the device would be authorized to use (i.e. playback) the MP3 data; if it was a "0," the device would refuse to play the music.

In addition to pre-authorization for certain classes of devices, the usage control string can also include bits indicating the number of permitted playbacks. This data can be encoded in bits seven through nine, representing eight possibilities:

0 – no playback permitted

5 1 – single playback permitted

2 – two playbacks permitted

3 – three playbacks permitted

4 – four playbacks permitted

5 – five playbacks permitted

10 6 – 10 playbacks permitted

7 – unlimited playbacks permitted

8 – refer to associated data (within the watermark, or stored at a remote site) which specifies number of permitted playbacks.

15 The playback device may include a non-volatile store in which the number of permitted playbacks is stored for each track of music. The device would decrement this number at the beginning of each playback.

20 The usage control string can also include a two-bit field (bits ten and eleven) indicating recording permissions. A value of 0 means that data corresponding to the MP3 audio (regardless of digital format) should never be made available to another digital device. A value of 1 means that the data corresponding to the MP3 data may be made available once to another digital device. A value of 2 means that the data may be made available an unlimited number of times to other digital devices. (Value 3 is reserved.)

25 Another data field that can be included in an audio watermark is a rating that indicates age-appropriateness. Music with violence or sexual themes might be given a rating akin to the MPAA "PG-13" or "R" rating. Audio appliances may be programmed to recognize the rating of incoming music, and to interrupt playback if the rating exceeds a certain threshold setting. Various known techniques can be employed to assure that such settings cannot readily be changed, e.g., by juvenile listeners.

30 Another data field that can be included in an audio watermark is a date field. This field can indicate either the date the music was watermarked, or a date in the future on which certain rights associated with the music should change. Some consumers, for example may not wish to

purchase perpetual playback rights to certain musical selections. The right to play a selection for 6 months may suffice for many consumers, especially if the price is discounted in view of the limited term. Such an arrangement would not be wholly disadvantageous to music distributors, since some consumers may end up purchasing music twice if their initial assessment of a musical selection's appeal was too short-sighted. (Naturally, the playback equipment would require a source of real-time clock data against which the date field in the watermark can be checked to ensure that the playback rights have not yet expired.)

Another of the data fields that can be included in an audio watermark specifies technical playback parameters. For example, the parameter can cause the playback appliance to apply a spectral equalization that favors bass frequencies, or treble frequencies, or mid-range frequencies, etc. Other pre-configured equalization arrangements can similarly be invoked responsive to watermark data. Likewise, the parameter can invoke special-effects provided by the playback appliance, e.g., echo effects, reverb, etc. (Again, such parameters are usually represented in an abbreviated, coded form, and are interpreted in accordance with instructions stored in a memory (either in the playback appliance, or linked thereto).)

The same data fields and principles can be applied to non-audio content. In video, for example, watermarked data can adaptively control the display monitor or playback parameters (e.g., color space) to enhance the viewing experience.

Music Asset Management/Commerce

The majority of domestic music piracy is not organized. Rather, it is a crime of opportunity and convenience. If the crime were made more difficult, the alternative of obtaining a copy through legitimate channels would be less onerous. Similarly, if the procedure for obtaining a copy through legitimate channels were simplified, the incentive for piracy would be reduced. Watermarking facilitates both – making the crime more difficult, and making legitimate music acquisition easier.

Consider, for example, the pricing of music in conventional record stores. A CD (compact disk) may cost \$15, but its sale may be driven by just one or two popular songs on the disk. To obtain these songs, the consumers must purchase the entire disk, with perhaps a dozen songs of no particular interest. This, in essence, is a tying arrangement that benefits the record

labels while prejudicing the consumers. Given these circumstances, and a ready opportunity to make copies, it is not surprising that customers sometimes make illicit copies.

5 One classic technique of avoiding purchase of a complete collection of music, when only one or two songs is desired, is to record the music off the radio. While of dubious legality, this technique was popular in the era of combined cassette/radio players. However, the desired music was sometimes difficult to encounter in a radio broadcast, and the quality was less than superb.

10 The combined cassette/radio player has now evolved into a general purpose computer with wide-ranging functionality, and other sophisticated devices. Music can be acquired off the web, and can be recorded in various forms (e.g. in a personal MP3 player, stored on a hard disk, stored on a writeable CD-ROM, played back and recorded on analog cassette, etc., etc.). The quality can be quite high, and the erratic broadcast time problems of radio broadcasts have been overcome by the web's on-demand delivery mechanisms. (Moreover, the music can be downloaded in faster-than-realtime, a further benefit over recording-off-the-air techniques.)

15 One hybrid between the new and old is a novel radio (e.g., for use in a car) that has a "capture" button on the front panel (or other form of user interface, e.g., a Capture icon on a GUI). If a user hears a song they want to record and keep, they press the Capture button while the song is playing. In response, the radio device decodes a watermark embedded in the music, and thereby knows the identity of the music. The radio then makes a wireless transmission identifying the user and the desired song. A local repeater network picks up the wireless signal and relays it (e.g. by wireless rebroadcast, by modem, or other communication medium) to a music clearinghouse. The clearinghouse charges the user a nominal fee (e.g. via a pre-arranged credit card), and queues the music for download to a predetermined location associated with the user.

25 In one embodiment, the predetermined location is the user's own computer. If a "live" IP address is known for the user's computer, the music can be transferred immediately. If the user's computer is only occasionally connected to the internet, the music can be stored at a web site (e.g. protected with a user-set password), and can be downloaded to the user's computer whenever it is convenient.

30 In other embodiments, the predetermined location is a personal music library maintained by the user. The library can take the form, e.g., of a hard-disk or semiconductor memory array in which the user customarily stores music. This storage device is adapted to provide music data to

5 The personal music library can have its own internet connection. Or it can be equipped with wireless capabilities, permitting it to receive digital music from wireless broadcasts (e.g. from the clearinghouse). In either case, the library can provide music to the user's playback devices by short-range wireless broadcast.

10

15

20

25

30

indexed in accordance with one or more fields from the decoded watermark, and the resulting textual data from the memory (e.g. song title and artist) is annunciated or displayed to the user.

5 In most applications, such memory will require frequent updating. The RF receiver provides a ready mechanism for providing such updated data. In one embodiment, the radio "awakens" briefly at otherwise idle moments and tunes to a predetermined frequency at which updated data for the memory is broadcast, either in a baseband broadcast channel, or in an ancillary (e.g. SCA) channel.

10 In variants of the foregoing, internet delivery of updated memory data can be substituted for wireless delivery. For example, the artist/song title memory in the personal player can be updated by placing the player in a "nest" every evening. The nest (which may be integrated with a battery charger for the appliance) can have an internet connection, and can exchange data with the personal device by infrared, inductive, or other proximity-coupling technologies, or through metal contacts. Each evening, the nest can receive an updated collection of artists/song titles, and can re-write the memory in the personal device accordingly. By such arrangement, the watermark data can always be properly interpreted for presentation to the user.

15 The "Capture" concepts noted above can be extended to other functions as well. One is akin to forwarding of email. If a consumer hears a song that another friend would enjoy, the listener can send a copy of the song to the friend. This instruction can be issued by pressing a "Send" button, or by invoking a similar function on a graphical (or voice- or gesture-responsive) user interface. In response, the appliance so-instructed can query the person as to the recipient. The person can designate the desired recipient(s) by typing in a name, or a portion thereof sufficient to uniquely identify the recipient. Or more typically, the person can speak the recipient's name. As is conventional with hands-free vehicle cell phones, a voice recognition unit can listen to the spoken instructions and identify the desired recipient. An "address book"-like feature has the requisite information for the recipient (e.g., the web site, IP address, or other data identifying the location to which music for that recipient should be stored or queued, the format in which the music should be delivered, etc.) stored therein. In response to such command, the appliance dispatches instructions to the clearinghouse, including an authorization to debit the sender's credit card for the music charge. Again, the clearinghouse attends to delivery of the music in a desired manner to the specified recipient.

20

25

30

Still further, a listener may query the appliance (by voice, GUI or physical button, textual, gesture, or other input) to identify CDs on which the then-playing selection is recorded. Or the listener may query the appliance for the then-playing artist's concert schedule. Again, the appliance can contact a remote database, relay the query, and forward data from the watermark payload identifying the artist and/or song title to which the query relates. The database locates the requested data, and relays same back to the appliance for presentation (via a display, by machine speech, or other output) to the user. If desired, the user can continue the dialog with a further instruction, e.g., to buy one of the CDs on which the then-playing song is included. Again, this instruction may be entered by voice, GUI, etc., and dispatched from the appliance to the clearinghouse, which can then complete the transaction in accordance with pre-stored information (e.g. credit card account number, mailing address, etc.). A confirming message is relayed to the appliance for presentation to the user.

While the foregoing transactions require a link to a remote site or database, other watermark-based consumer services can be provided without such a link. For example, a user can query the appliance as to the artist or song-title of the selection currently playing. The appliance can consult the embedded watermark data (and optionally consult a memory to determine the textual names associated with coded watermark data), and provide the requested information to the user (e.g., by a display, annunciation, or other output).

The foregoing concepts (e.g. Capture, Send, etc.) can also be employed in connection with internet- rather than radio-delivery of music. (The following discussion is illustrated with reference to the "Capture" function, but it will be recognized that the other earlier-discussed features can be similarly implemented.)

There are many commercial web sites that sell audio (in CD form or otherwise), and offer limited free music downloads, (or music clips) as an enticement to lure consumers. But there are also a great number of music web sites that have no commercial pretense. They are hosted by music lovers strictly for the enjoyment of other music lovers. When music is downloaded from such a web site, the end-user's computer can analyze the digital data to decode watermark data therefrom. Again, the user can be presented with a "Capture" button that initiates a commercial transaction, by which a complete copy of the then-downloaded audio is sent to a prearranged storage location, and the user's credit card is debited accordingly. This transaction can occur

independently of the site from which the music is downloaded (e.g. through the clearinghouse referenced above).

While the "Capture" button can be presented on the web-site, this would generally not be in keeping with the non-commercial nature of such web sites. Instead, in an exemplary embodiment, the Capture feature is a software program resident at the user's computer. When this software program is invoked by the user, a socket channel is instantiated between the user's computer and the clearinghouse over the then-existing internet connection. The decoded watermark data and user ID is transmitted to the clearinghouse over this channel, without interrupting the user's other activity (e.g. downloading music from the non-commercial web site). In response, the clearinghouse transmits the music to the prearranged location and attends to billing.

In some embodiments, a watermark detector is included as part of the operating system, and constantly monitors all TCP/IP, or other internet, data received by the user's computer, for the presence of watermarks. In such case, when the Capture feature is invoked, the program examines a memory location in which the operating system stores the most-recently received watermark data. In another embodiment, the computer does not monitor all internet traffic for embedded watermark data, but includes an API that can be called by the Capture program to decode a watermark from the data then being received. The API returns the decoded watermark data to the Capture program, which relays same to the clearinghouse, as above. In still another embodiment, the watermark decoder forms part of the Capture program, which both decodes the watermark and relays it to the clearinghouse when the Capture program is invoked by the user.

There are various techniques by which the Capture program can be selectively invoked. One is by a keyboard macro (e.g. by a combination of keyboard keys). Another is by a program icon that is always presented on the screen, and can be double-clicked to activate. (Again, confirmation processes may be called for, depending on the likelihood of inadvertent invocation.) Many other techniques are likewise possible.

In the just-contemplated scenario, the Capture operation is invoked while the user is downloading music from a non-commercial web site. This seems somewhat redundant, since the downloading -- itself -- is transferring music to the user's computer. However, the Capture operation provides added value.

In the case of streaming audio, the audio is not typically stored in a location in which it can be re-used by the consumer. It can be listened-to as delivered, but is then gone. Capturing the audio provides the user a copy that can be played repeatedly.

5 In the case of downloaded music files, the music may have been encoded to prevent its recordal on other devices. Thus, while the user may download the music onto a desktop computer, copy-prevention mechanisms may prevent use of that file anywhere else, e.g. on a portable music appliance. Again, Capturing the audio provides the user a copy that can be transferred to another device. (The music file provided by the clearinghouse can have copy-prevention limits of its own – e.g., the file can be copied, but only once, or the file can be copied
10 only onto devices owned by the user.)

(Confirmation of device ownership can be implemented in various ways. One is to identify to the clearinghouse all music devices owned by a user at the time the user registers with the clearinghouse (supplemented as necessary by later equipment acquisitions). Device IDs associated with a user can be stored in a database at the clearinghouse, and these can be encoded
15 into the downloaded music as permitted devices to which the file can be copied, or on which it can be played.)

The commerce opportunity presented by non-commercial music web-sites is but one enabled by digital watermarks. There are many others.

To take one example, consider the media by which music and artists are presently
20 promoted. In addition to radio airtime, these include music videos (a la MTV), fan magazines, web advertisements, graphical icons (e.g. the Grateful Dead dancing bears), posters, live events, movies, etc. Watermarked data can be used in all such media as a link in a commercial transaction.

A poster, for example, typically includes a photo of the artist, and may comprise cover
25 art from a CD. The photo/art can be digitally watermarked with various types of data, e.g., the artist's identify, the record label that distributes the artist's work, the music project being particularly promoted by the poster (e.g. a CD, or a concert tour), a fan web-site related to the artist, a web-site hosted by the record label for selling audio in CD or electronic form, a web-site from which free music by the artist can be downloaded, data identifying the poster itself, etc.

A user, equipped with a portable appliance that merges the functions of palmtop computer and digital camera, can snap an image of the poster. The processor can decode the watermarked data, and initiate any of various links based on the decoded data.

In an exemplary embodiment, after snapping the picture, the user invokes a software program on the device that exposes the various links gleaned from the snapped image data. Such a program can, for example, present the option of linking to the artist's fan web site, or downloading free streaming audio or music clips, or ordering the promoted CD, or requesting the above-noted clearinghouse to download a personal copy of selected song(s) by the artist to the user's personal music library, etc. (The device is presumed to have a wireless internet link. In devices not having this capability, the requested actions can be queued and automatically executed when a link to the internet is available.)

Still more complex transactions can be realized with the use of a remote database indexed by digital watermark fields decoded from the poster. For example, the poster may promote a concert tour. Fields of the digital watermark may identify the artist (by a code or full text), and a web site or IP address. The user appliance establishes a link to the specified site, and provides the artist identifier. In response, the site downloads the tour schedule for that artist, for display on the device. Additionally, the downloaded/displayed information can include a telephone number that can be used to order tickets or, more directly, can indicate the class of seats still available at each (or a selected) venue, and solicit a ticket order from the user over the device. The user can supply requested information (e.g. mailing address and charge card number) over the return channel link (wireless or wired, as the case may be), and the ticket(s) will be dispatched to the user. In the case of a wireless link, all of this can occur while the user is standing in front of the movie poster.

Similar systems can be implemented based on watermark data encoded in any other promotional media. Consider music videos. Using known TV/computer appliances, watermark data added to such videos can readily be decoded, and used to establish links to audio download, CD-sales, fan club, concert ticket outlet web sites, etc., as above.

Even live events offer such watermark-based opportunities. The analog audio fed to public address or concert speakers can be watermarked (typically before amplification) to encode plural-bit digital data therein. A next generation personal music appliance (e.g. one with a wireless interface to the internet) can include analog record capability (e.g. a built-in

microphone, analog-to-digital converter, MP3 encoder, coupled to the unit's semiconductor memory). A user who attends a live event may record an excerpt of the music. The watermark can then be decoded, and the extracted data used to access the links and commerce opportunities reviewed above.

5 Cinema movies offer both audio and visual opportunities for watermark-based commerce opportunities. Either medium can be encoded to convey information of the types reviewed above. A personal appliance with image- or audio-capture capabilities can capture an excerpt of the audio or imagery, decode the watermark data therefrom, and perform any of the linking, etc., functions reviewed above.

10 The consumer-interest watermarks reviewed above are only exemplary. Many others will be recognized as useful. For example, promotional clips presented before a feature film presentation can include watermark data that point (either by a literally encoded web address link, or by an ID code that indexes a literal link in a remote link database) to reviewer critiques of the previewed movies. Watermark data in a featured film presentation can lead to web sites
15 with information about the movie stars, the director, the producer, and can list other movies by each of these persons. Other watermark-conveyed web links can present opportunities to buy the movie on videotape, to purchase the movie soundtrack, to buy movie-related toys and games, etc.

More on Device Control

20 Much of the foregoing has focused on watermark encoding to provide enhanced customer experiences or opportunities. Naturally, watermarks data can alternatively, or additionally, serve the interests of the media owner.

To illustrate, consider watermarked music. The media owner would be best served if the watermark serves dual purposes: permissive and restrictive. Permissively, music appliances can
25 be designed to play (or record) only music that includes an embedded watermark signaling that such activity is authorized. By this arrangement, if music is obtained from an unauthorized source and does not include the necessary watermark, the appliance will recognize that it does not have permission to use the music, so will refuse requests to play (or record).

30 As noted, music appliances can respond restrictively to the embedded watermark data to set limits on use of the music. Fields in the watermark can specify any or all of (or others in addition to) (a) the types of devices on which the music can be played (b) the types of devices on

which the music can be recorded; (c) the number of times the music can be played; (d) the number of times the music can be recorded, etc.

The device restrictions (a) and (b) can be of various types. In some embodiments, the restrictions can identify particular units (e.g. by serial number, registered owner, etc.) that are authorized to play/record the encoded music. Or the restrictions can identify particular classes of units (e.g., battery-powered portable players with music memories of less than 50 megabytes, disk-based dedicated music appliances, general purpose personal computers, etc.) Or the restrictions can identify particular performance quality criteria (e.g., two channel, 16-bit audio at 44.1KHz sample rate, or lower quality).

The use restrictions (c) and (d) can likewise be of various types. Examples include "do not copy," "copy once only," "unrestricted copying permitted," "play once," "play N times" (where N is a parameter specified elsewhere in the watermarked data, or by reference to a database indexed by a watermark data field), "unrestricted playing permitted," etc.

It is straightforward to design a music appliance to respond to usage limits of zero (e.g. "do not copy") and infinity (e.g. "unrestricted copying permitted," and "unrestricted playing permitted"). The device simply examines one or more bits in the watermark data, and permits (or refuses) an operation based on the value thereof.

Implementation of the other usage-control restrictions can proceed in various ways. Generally speaking, the stored music can be altered to give effect to the usage-control restrictions. For example, if the music is "record-once," then at the time of recording, the appliance can alter the music in a fashion indicating that it now has "do not record" status. This alteration can be done, e.g., by changing the watermark data embedded in the stored music (or adding watermark data), by changing other data stored in association with the music, etc. If the original signal is stored (as opposed, e.g., to a streaming signal, such as an internet or wireless transmission), it too should be so-altered.

Likewise with playback limitations. The number of playbacks remaining can, e.g., be encoded in an updated watermark in the music, be tracked in a separate counter, etc.

More particularly considering the "copy once" usage restriction, an illustrative embodiment provides two distinct watermark payload bits: a "copy once" bit and a "copy never" bit. When originally distributed (whether by internet, wireless, or otherwise), the "copy once" bit is set, and the "copy never" bit is un-set.

When music encoded in this fashion is provided to a compliant recording device, the device is authorized to make one copy. (A compliant device is one that recognizes encoded watermark data, and behaves as dictated by the watermark.) When this privilege is exercised, the recording device must alter the data to ensure that no further copying is possible. In the illustrated embodiment, this alteration is effected by the recording device adding a second watermark to both the music, with the "copy never" bit asserted. The second watermark must generally be encoded in an "orthogonal" domain, so that it will be detectable notwithstanding the continued presence of the original watermark. Compliant equipment must then check for both watermarks, and refuse to copy if either is found to have the "copy never" bit asserted.

One advantage to this arrangement is that if the watermark signal has undergone some form of corruption (e.g. scaling or resampling), the first watermark may have been weakened. In contrast, the second watermark will be native to the corrupted signal, and thus be more easily detected. (The corruption may also contribute to the orthogonality of one watermark relative to the other, since the two watermarks may not have precisely the same time base or other foundation.)

An alternative approach is not to encode the "copy never" bit in the original music, but leave this bit (in whatever manifestation) blank (i.e. neither "1" nor "0"). In transform-based watermark techniques, this can mean leaving transform coefficient(s) corresponding to the "copy never" bit un-changed. If the watermarking is effected in the temporal sample domain (or spatial domain, for image data), this can mean leaving certain samples (pixels) unmodified. The recording device can then alter the transform coefficients and/or samples as necessary to assert the previously-unencoded "copy never" bit when the permitted recording is made.

In such a system, compliant recording devices check for the "copy never" bit in the sole watermark, and refuse to make a copy if it is asserted (ignoring the value of any "copy once" bit).

A third approach to "copy once" is to set both the "copy once" and "copy never" bits, but set the former bit very weakly (e.g. using lower gain and/or high frequency DCT coefficients that do not survive certain processing). The frail "copy once" bit is designed not to survive common corruptions, e.g., resampling scaling, digital to analog conversion, etc. To further assure that the "copy once" bit is lost, the recording device can deliberately add a weak noise signal that masks

this bit (e.g. by adding a noise signal in the frequency band whose DCT coefficient conveys the "copy once" bit). In contrast, the "never copy" bit is unchanged and reliably detectable.

In such a system, compliant devices check for the "copy once" bit in the sole watermark, and refuse to make a copy if it is not detected as set.

5 These three examples are but illustrations of many possible techniques for changing the rights associated with a work. Many other techniques are known. See, e.g., the proposals for watermark-based copy control systems for digital video at the Copy Protection Technical Working Group, <http://www.dvcc.com/dhsg/>, from which certain of the foregoing examples are drawn. See also Bloom et al, "Copy Protection for DVD Video," IEEE Proceedings, Special
10 Issue on Identification and Protection of Multimedia Information, June, 1999.

Scaleability

15 One feature that is desirable in many detector embodiments is scaleability. This refers to the ability of a detector to scale its computational demands to match the computational resources available to it. If a detector is running on a high performance Pentium III workstation, it should be "doing more" than if the same detector is running on a slow microcontroller. One way
20 scalability can be achieved is by processing more or less chunks of input data (e.g. temporal excerpts of music, or blocks/macroblocks of pixels in a frame of video data) to decode watermarks. For example, an input audio stream might be broken into chunks of one second each. A fast processor may complete decoding of each chunk in less than a second, permitting it
25 successively to process each chunk in the data stream. In contrast, a slow processor may require two and a half seconds to decode the watermark from a chunk. While it is processing a first chunk, the second and third pass by un-decoded. The processor next grabs and processes the fourth chunk, permitting the fifth and sixth to pass by un-encoded.

25 The detector running on the fast processor is clearly more difficult to "fool," and yields a decoded watermark of higher confidence. But both systems decode the watermark, and both operate in "real time."

30 The skipping of input data in the temporal (e.g. music or video) or spatial (e.g. image or video) domain is but one example of how scaleability can be achieved. Many other approaches are known to those skilled in the art. Some of these alternatives rely on spending more or less time in the data analysis phases of watermark decoding, such as cross-correlation operations.

Reference has been made to watermarked UIDs as referring to a database from which larger data strings (e.g. web addresses, musician names, etc.) can be retrieved. In some embodiments, the data record referenced by a UID can, in turn, point to several other database records. By such arrangements, it is often possible to reduce the payload of the watermark, since a single UID reference can lead to several different data records.

Production Tools

In the prior art, the watermark embedded in a source material is typically consistent and static through a work – unchanging from beginning to end. But as will be recognized from the foregoing, there are many applications that are better served by changing the watermark data dynamically during the course of the work. According to another aspect of the invention, a production tool is provided that facilitates the selection and embedding of dynamically-changing watermark data. One such embodiment is a software program having a user interface that graphically displays the different watermark fields that are being embedded in a work, and presents a library of data (textually or by icons) that can be inserted into each field, and/or permits the user to type in data to be encoded. Another control on the UI controls the advance and rewind of the media, permitting the user to determine the location at which different watermark data begins and ends. Graphical paradigms known from video- and audio-editing tools can be used to indicate the starting and ending frames/samples for each different watermark payload.

Such a tool can be of the standalone variety, or can be integrated into the desktop audio- and video- production and editing tools offered by vendors such as Avid, Adobe, Jaleo, Pinnacle Systems, SoundForge, Sonic Foundry, Xing Technology, Prosoniq, and Sonic Desktop Software.

Payment-Based Systems

Another aspect of the present invention is the use of anonymous payment tokens that can be used to obtain content on the web. In one embodiment, a token comprises a 128-bit pseudo-random number, to which additional bits identifying an issuing bank (or other issuing institution) are appended. (The additional bits can be the IP address of a web server of the bank, a routing number identifying the bank for electronic wire transfers, or other identifier.) The 128-bit

numbers are randomly generated by the bank – commonly as needed – and each represents a fixed increment of money, e.g. ten cents.

A consumer wishing to have a store of currency for such commerce pays the bank, e.g., \$10 in exchange for 100 tokens. These tokens are transferred electronically to disk or other storage in the consumer's computer in response, e.g., to a credit card authorization, or may be provided by diskette or other storage medium over the counter at a bank branch (in which case the consumer thereafter copies the numbers into storage of his or her computer). (Outlets other than banks can of course be employed for distributing such numbers, much in the manner that convenience and many grocery stores commonly issue money orders.)

Imagine that the consumer wishes to view the final quarter of a Trailblazer basketball game that aired on television a week ago. (The consumer may have either missed the game, or may have seen it but wants to see the last quarter again.) The user directs a web browser to a web site maintained for such purpose and performs a search to identify the desired program. (Typically, the web site is maintained by the proprietor that holds the copyright in the material, but this need not be the case. Some material may be available at several web sites, e.g., maintained by ABC Sports, the National Basketball Association, and Sports Illustrated.) The search can use any of various known search engines, e.g., Infoseek, Verity, etc., and can permit searching by title terms, keywords, date of airing, copyright owner, etc. By typing in, e.g., the keyword 'Trailblazers' and the date '4/26/99,' the consumer is presented a listing of videos available for download. One, hopefully, is the requested game. With each listing is an indication of an associated nominal charge (e.g. 80 cents).

On clicking on a hypertext link associated with the desired basketball game, the viewer is presented a further screen with one or more options. The first of the listed options is the entire game, with commercials. The charge is the nominal charge presented on the earlier screen (i.e. 80 cents). Other options may include the first, second, third, and fourth quarters of the game individually, each of which – save the last, costs 20 cents. The last may be charged at a premium rate, e.g., 30 cents. Clicking on the desired video option yields a further screen through which payment is effected.

To pay for the requested video, the consumer instructs his or her computer to transfer three of the earlier-purchased tokens over the web to the video provider. Various user interface metaphors can be employed to facilitate this transfer, e.g., permitting the user to type the amount

of money to be transferred in a dialog box presented on-screen, or dropping/dragging icons representing tokens from an on-screen "wallet" to an on-screen "ticket booth" (or over an icon or thumbnail representing the desired content), clicking on an "increment" counter displayed adjacent the listing of the content, etc. Once the consumer has authorized a transfer of sufficient tokens, the consumer's computer sends to the web site (or to such other web address as HTML encoding in the viewed web page may indicate) the tokens. This transmission simply takes the form of the three 128+ bit numbers (the '+' indicating the bank identifier) – in whatever packet or other format may be used by the internet link. Once dispatched in this manner, the tokens are deleted from the user's computer, or simply marked as spent. (Of course, in other embodiments, a record of the expenditure may be stored in the consumer's computer, e.g., with the token contents and a record of the audio or video purchase to which they were applied.)

Since the amount of money is nominal, no encryption is provided in this embodiment, although encryption can naturally be provided in other embodiments (e.g., either in sending the tokens from the user to the web site, or earlier, in sending the tokens to the user). As will be seen, provided that the media provider immediately sends the tokens to the bank in real time, encryption is a nice feature but not mandatory

On receipt of the token data, the web site immediately routes the token data to the identified bank, together with an identifier of the media provider or account to which the funds represented thereby are to be credited. The bank checks whether the 128-bit numbers have been issued by that bank, and whether they have already been spent. If the numbers are valid, the bank updates its disk-based records to indicate that the three tokens have been spent and that the bank now owes the media supplier 30 cents, which it may either pay immediately (e.g., by crediting to an account identified by the media provider) or as one lump sum at the end of the month. The bank then sends a message to the web site confirming that the tokens were valid and credited to the requested account. (Optionally, a message can be sent to the purchaser of the tokens (if known), reporting that the tokens have been redeemed.)

In response, the web site begins delivery of the requested video to the consumer. In the illustrated embodiment, the video is watermarked prior to delivery, but otherwise sent in unencrypted fashion, typically in streaming format, but optionally in file format. (Encryption can be used in other embodiments.) The watermarking in the illustrated embodiment is accomplished

on-the-fly and can include various data, including the date of downloading, the download site, the destination IP address, the identity of the purchaser (if known), etc.

The large size of the video and the small charge assessed therefor provide disincentives for the consumer making illicit copies. (Especially as to archival material whose value decays with time, there is not much after-market demand that could be served by illicit copies, making third party compilation of such material for re-distribution financially unattractive. First run video, and material that keeps a high value over time, would not be as well suited for such distribution, and could better employ technology disclosed elsewhere herein.)

In some embodiments, the integrity of the received video is checked on receipt. This feature is described below in the section entitled Watermark-Based Receipts.

In the illustrative system, nothing in the tokens indicates the identity of the purchaser. The web site knows the IP address of the site to which video was delivered, but need not otherwise know the identity of the purchaser. The bank would probably maintain a record of who purchased the tokens, but need not. In any event, such tokens could thereafter be exchanged among consumers, resulting in anonymity from the bank, if desired.

As described above, the video excerpts from which the consumer can select include commercials. At some sites, video may be provided from which the commercials have been excised, or which is delivered in a manner that skips past the commercials without transmitting same to the consumer. Such video will naturally command a premium price. In some embodiments, the difference in price is electronically credited as compensation to accounts maintained for (or by) the advertisers, whose advertisements are not being viewed by such consumers. (The identification of advertisers to be credited is desirably permanently encoded in the video, either throughout the video (if the video has had the commercials removed therefrom), or by data in the commercials themselves (which commercials are skipped for transmission to the consumer, but can still be decoded at the video head-end. Such encoding can be by in-band watermarking or otherwise.)

While the foregoing discussion particularly considered video as the desired content, the same principles are equally applicable in connection with audio, still imagery, and other content.

The token-based payment method is but one of many that can be employed; the literature relating to on-line payment mechanisms is extensive, and all such systems can generally be here-employed.

Tracking 128-bit tokens can be a logistical problem for the bank. One approach is to have a memory with 10^{128} locations, and at each location store a two-bit value (e.g. 00=never issued; 01=issued but not spent; 10=issued and spent; 11=reserved). More complete data could alternatively be stored, but such a memory would be impractically large.

5 One alternative approach is to hash each 128-bit number, when issued, to a much smaller key value (e.g. 20 bits). A memory with 10^{20} locations can be indexed by this key. Each such location can include four data: an issued 128-bit token number that hashes to that value, first and second date fields indicating the date/time on which that token was issued and redeemed, respectively, and a link specifying the address of a next memory location. That next memory
10 location (outside of the original 10^{20} locations) can include four more data, this time for a second issued-128-bit token number that hashed to the original key value, two date fields, and again with a link to a subsequent storage location, etc.

15 When a 128-bit random number is generated, the original memory location indexed by the hash code of that number is checked for an earlier number of the identical value (to avoid issuance of duplicate tokens). Each successive location in the linked chain of memory locations is checked for the same 128-bit number. When the end of the linked chain is reached, the bank knows that the 128-bit random number has not previously been issued, and writes that number in the last-addressed location, together with the date of issuance, and a link to a next storage location.

20 When a 128-bit token is received, the same linked-list processing occurs to identify a first location, and to thereafter step through each subsequent location until a match is found between the token number and the number stored in one of the linked memory locations. When found, that number is marked as redeemed by writing a redemption date/time in the corresponding field. If the search reaches the end of the linked chain without finding a match between the stored
25 numbers and the token number, the token is treated as invalid (i.e. not issued by that bank).

30 Other manners of tracking the large number of possible token numbers can of course be used; the foregoing is just exemplary. Or the tokens needn't be tracked at all. Such an arrangement is highly practical if the token has sufficient bits. With the illustrated 128 bits, for example, the chance of two identical tokens being issued is infinitesimally small, so checking for duplicate issuance can be omitted if desired. In such case, the bank can simply maintain an ordered list of the token numbers still outstanding and valid. As new tokens are dispensed, their

token numbers are added to the list. As tokens are redeemed, their numbers are deleted from the list. Known list processing techniques can be employed to speed such search, update, and delete actions.

5 Watermark-Based Receipts

Pay-for-content applications commonly assume that if content is transmitted from a server (or head-end, etc.), it is necessarily received. Sometimes this assumption is wrong. Network outages and interruptions and internet traffic load can diminish (e.g., dropped video frames), or even negate (e.g., failed delivery), expected consumer enjoyment of content. In such cases, the consumer is left to haggle with the content provider in order to obtain an adjustment, or refund, of assessed charges.

Watermarks provide a mechanism for confirming receipt of content. If a watermark is detected continuously during a download or other delivery event, a software program (or hardware device) can issue an electronic receipt attesting that the content was properly delivered. This receipt can be stored, and/or sent to the content distributor to confirm delivery.

In one embodiment, a content receiving device (e.g., computer, television or set-top box, audio appliance, etc.) periodically decodes a watermark from the received content to confirm its continued reception. For example, every five seconds a watermark detector can decode the watermark and make a record of the decoded data (or simply record the fact of continued detection of the same watermark). When a changed watermark is detected (i.e., reception of a different content object begins), the duration of the previously-received content is logged, and a receipt is issued.

In a related embodiment, the last portion (e.g., 5 seconds, frame, etc.) of the content bears a different "end of content" watermark that triggers issuance of a receipt. Such a watermark can indicate the length of the content, to serve as a cross-check against the periodic watermark polling. (E.g., if periodic sampling at 2 second intervals yields 545 samples corresponding to the same content, and if the "end of content" watermark indicates that the content was 1090 seconds long, then receipt of the entire content can be confirmed.)

In another embodiment, the watermark can change during the course of the content by including, e.g., a datum that increments every frame or other increment of time (e.g., frame number, time stamp, etc.). A watermark detector can monitor the continued incrementing of this

datum throughout the content to confirm that no part was garbled (which would destroy the watermark) or was otherwise missing. Again, at the end of delivery, the receiving system can issue a confirmation that XXX frames/seconds/etc. of the identified content were received.

5 One application of such technology is to bill for content based on receipt, rather than transmission. Moreover, billings can be adjusted based on percentage of content-value received. If delivery is interrupted mid-way through (e.g., by the consumer disabling the content-receiving device), the nominal billing for the content can be halved. Some prolonged content, e.g., televised/web-broadcast university classes, cannot be "consumed" in one session, and are thus particularly well suited for such pay-as-you-consume billing.

10 Another application of such technology is in advertising verification. Presently, ads are tracked by transmission or, less frequently, by detection of an embedded code on receipt (*c.f.*, Nielsen Media Research's patents 5,850,249 and 5,737,025). However, such reception-detectors – once triggered – generally do not further note the length of time that the advertising was received, so the same data is produced regardless of whether only five or fifty seconds of a
15 commercial is presented. Watermark monitoring as contemplated herein allows the duration of the advertising impression to be precisely tracked.

In one application of this technology, recipients of advertising are provided incentives for viewing advertising in its entirety. For example, a content-receiving device can include a watermark detector that issues a receipt for each advertisement that is heard/viewed in its
20 entirety. These receipts can be redeemed, e.g., for content tokens as described elsewhere herein, for monetary value, etc. In some embodiments, receipts are generic and can all be applied to a desired premium, regardless of the advertisements through which they were earned. In other embodiments, the receipts are associated with the particular advertisers (or class of advertisers). Thus, a TV viewer who accumulates 50 receipts from advertising originating from Procter &
25 Gamble may be able to redeem same for a coupon good for \$2.50 off any Procter & Gamble product, or receipts from Delta Airlines may be redeemed for Delta frequency flier miles (e.g., at a rate of one mile per minute of advertising). Such incentives are particularly useful in new forms of media that give the consumer enhanced opportunities to fast-forward or otherwise skip advertising.

30 (Although the foregoing "receipt" concept has been described in conjunction with watermark data (and use of watermark technology is believed to be inherently advantageous in

this application), the same principles can likewise be implemented with ancillary data conveyed by other means.)

Master Global Address

5 As suggested above, it is desirable that each piece of content have a web address (the "Master Global Address" (MGA), or "Master IP Address") associated with it. Such address is typically conveyed with the content, e.g., by an IP address watermarked therein.

10 Consider a consumer who downloads a streaming video having an English language soundtrack. The viewer may not speak English, or may otherwise prefer to listen to the soundtrack in another language. The user can decode the watermark data embedded in the video and initiate a link to the associated web address. There the user is presented with a list of soundtracks for that content object in other languages. The viewer can click on the desired language and receive same via a second simultaneous transmission (e.g., a second socket channel). The consumer's audio/video appliance can substitute the desired audio track for the
15 default English track.

If the streaming video and the alternative soundtrack are hosted on the same server, synchronization is straightforward. The process governing transmission of the alternative soundtrack identifies the process that is streaming video to the same IP address. Based on SMPTE, or other time/frame data, the former process syncs to the latter. (If the two data streams
20 don't originate through the same server, time/frame data can be relayed as necessary to the alternative soundtrack server to effect synchronization.)

Another application of the Master Global Address is to serve as a point to which monitoring stations can report the presence, or passage, of content. Consider, for example, a copyright-aware node through which content signals pass, e.g., a computer node on a network, a
25 satellite transponder, etc. Whenever the node detects passage of a media object (e.g., by reference to a file extension, such as MP3, JPG, AVI, etc.), it sends a "ping" over the internet to the address encoded in the object, simply reporting passage of the object. Similar monitoring facilities can be provided in end user computers, e.g., reporting FileOpen, FileSave, Printing, or other use of content bearing MGA data.

30 This system can be expanded to include "ping" and "pong" phases of operation. When a software application (or a user appliance, such as a video or audio playback device) encounters a

media object (e.g., at time of file open, at time of playback, etc.), it pings the MGA site to report the encounter. The MGA site "pongs" back, responding with instructions appropriate to the encounter. For example, if the object requires payment of a fee before full functionality or access is to be granted, the MGA site can respond to the application with instructions that the object be used (e.g., played back) only in some crippled state preventing the user's full enjoyment (e.g., impaired resolution, or impaired sound quality, or excerpts only, etc.). The MGA site can also inform the user application of the terms (e.g., payment) by which full functionality can be obtained. The application can graphically or audibly present such information to the user, who can authorize a payment, if desired, so that the content can be enjoyed in a less- (or un-) crippled state. On receipt of the payment authorization, the MGA site can inform the user application that enhanced access/usage rights have been purchased, and that the application may proceed accordingly.

Yet another application of the MGA is to present the user of a content object a menu of options that is customized to that object.

In current graphical operating systems, when a user clicks on an icon (e.g., with the right mouse button), a menu is presented detailing actions that can be undertaken in connection with the icon, or the file represented thereby. Such options are pre-programmed (i.e., static), and are typically determined by the operating system based solely on the file extension.

In accordance with this aspect of the present invention, clicking on an icon representing a media object initiates an internet link to the MGA site associated with the object. The MGA site responds with data that is used to customize the menu of options presented to the user in connection with that particular object.

Consider an icon representing a JPG image file. Right-clicking on the icon may yield a menu that gives the user various options presented by the operating system (e.g., delete, compress, rename), and additional options customized in accordance with data from the object's MGA site. These customized options may include, e.g.,

- (a) open in 100x150 pixel format for free;
- (b) open in 480x640 pixel format for ten cents;
- (c) open in 960x1280 pixel format for twenty cents;
- (d) purchase rights to use this image in a newsletter having a circulation of under 1000 for \$1.25;

(e) display a complete listing of license options.

Clicking on options (b) or (c) initiates a commerce application through which funds are electronically transferred to the MGA site (by the above-described tokens or otherwise). In response, the MGA site responds (e.g., with TCP/IP or HTML instructions) authorizing an application on the user's computer to open the file in the requested manner. (The default application for JPG applications can then automatically be launched, or the computer may first query the user whether another application should be used instead.)

Clicking on option (d) proceeds as above, and permits full use of the image on the computer. Moreover, the MGA site sends a digital certificate to the user's computer memorializing the usage rights purchased by the consumer.

In this particular arrangement, no access control is placed on the content, e.g., by encryption, secure container technology, or the like. The nominal fees, and the ease of licensing, make it simple for the user to "do the right thing" and avoid copyright liability. In other embodiments, of course, known access control techniques can be used to limit use of the object until the requisite payment has been made.

Naturally, records of all such transactions are also logged at the MGA site.

Clicking on option (e) opens a browser window on the user's computer to a web site that presents a complete listing of license options available for that image. (The address of this web site is included in customization data relayed to the user device from the MGA site, but not explicitly shown to the user on the menu.) Through such web site, the user can select desired rights, effect payment, and receive the necessary authorization for software applications on the user's computer (or other media appliance) to open and/or process the content.

The object on which the user "clicks" needn't be an icon. It can be an image or other graphical representation. (And a "click" isn't necessary; a voice command or other signal may be used to the same effect with an audio clip or selection.)

Consider the popular merchandising of books and CDs over the internet. A JPG or other image file depicting the cover of a book, or the artwork of a CD cover, can be treated as a media object, and can include a watermarked MGA pointer. Right-clicking on such an image of a book cover could, through the MGA site, present to the user a menu of options that includes – in addition to those normally presented in conjunction with a JPG file – the following:

- (a) "See the review of this book published in the New York Times on April 19, 1999"
- (b) "See the list of reviews of this book at Amazon.com"
- (c) "Enter your own review of this book, for posting on Amazon.com"
- (d) "See today's sales rank of this book at Amazon.com"
- 5 (e) "Purchase this book from Amazon.com for \$16.95"
- (f) "Purchase this book from Barnesandnoble.com for \$19.95 and receive a \$5.00 credit towards your next purchase"
- (g) "Link to the web site that tells about the release of this title as a motion picture (presently scheduled to open on October 10, 1999)"
- 10 (h) "Link to the Yahoo listing of web sites relating to this book"
- (i) "Search Lycos for listings relating to this book."

If the user selects one of the purchase options from the menu, a pre-stored e-commerce profile -- containing the user name, credit card number, billing address, ship-to address, etc., possibly in the form of an encrypted object -- could be sent to the MGA site (or to the bookseller) to effect the purchase, or such selection could initiate display of additional screens or sub-menus through which the user would manually enter or select such information for transmission.

15

Others of the selections cause a new browser window to open on the user's computer, opening to a URL specified in data relayed from the MGA site but not displayed to the user in the menu. Appropriate HTML instructions can be generated to effect a particular query or other operation at the specified URL.

20

In some embodiments, the customized menu presents only a single choice in addition to those normally provided by the operating system, e.g., "Link to home." Clicking on this option opens a browser window to a home page at the MGA for that object. On that page, the user is presented with all of the foregoing options, and more (possibly including advertising graphics or multi-media). Such objects can serve as powerful marketing agents. Returning to the example discussed above, a JPG image file of a book cover may have, as its MGA, a web page hosted by a particular bookseller, providing purchase options and other information for that book.

25

Marketing of books (or CDs, or cars, or consumer appliances, or virtually anything else) can be effected by disseminating such vendor-issued JPGs as widely as possible. Some book cover

30

JPGs may be distributed by Amazon.com, others by Barnes&Noble.com, others by Borders.com

– each pointing back to a different MGA through which purchase transactions for that book may be performed.

Returning to the MGA-customized menus, these needn't be limited to menus resulting from clicking on an icon or image (or signaling during an audio excerpt). Drop-down menus in application programs can likewise be populated with customized options, in accordance with customization data obtained from the MGA site for the object presently being accessed or used. Most graphical operating systems and application programs have well developed toolsets permitting such menu customization. Again, other data relayed from the MGA site is not shown to the user, but is employed by the computer (e.g., a browser program) to carry out menu options selected by the user.

Again the foregoing techniques are equally applicable for still images, audio, video, and other forms of content, and can readily be adapted for use both with general purpose computers, software applications, and specialized media appliances.

While, for expository convenience, the foregoing discussion contemplated embedding a literal URL address in the object as the MGA, more typically this is not the case. Instead, the MGA more commonly comprises identification data for the object (e.g. a 128-bit random ID), together with the URL for a name server computer that serves many (perhaps millions) of such objects (an example of the latter is the Digimarc MarcCentre server).

To obtain the desired data as detailed above, the user's computer (sometimes termed a client computer) links to the name server computer and provides the ID of the object being processed. The name server computer uses this ID to query a database, and obtains from the database the current IP address to which such queries should be routed. The name server computer can relay the request from the client computer to the correct destination address, or can return the correct destination address to the client computer, which can initiate such a link itself. By such arrangement, the IP address ultimately associated with an object can be easily changed as needed, simply by changing the corresponding record in the name server database, without rendering obsolete legacy objects having out-of-date addresses encoded therein.

In some embodiments, the URL of the name server needn't be included in the watermark. In the absence of a specified URL, the client computer may direct such links to a default name server address instead (stored locally or remotely). If that server doesn't recognize the object ID, it can return an error code, or pass the query on to other name servers. Those servers, in turn,

can pass the query along to still other name servers if they don't recognize the object ID. In this fashion, an exponentially-large number of name servers might be quickly polled for information relating to the identified object. Alternatively, rather than encoding the complete IP address of the name server in an object watermark, the first N (e.g., 16) bits of the object ID might be used as a short-hand for one of 65,536 predetermined name server addresses, in accordance with data stored locally (e.g., on RAM or disk in the user's computer) or remotely (e.g., at a default name server IP address).

While the basic concept idea behind embedding MGA data within an object is to point to a repository of data about the object, a pointer the other way may be achieved as well.

As noted, the "ping" application of MGA data permits an MGA site to be informed of sites through which its object passes. More generally, the MGA site can log the originating address of each query it receives. Each such address can be presumed to have (or have had) a copy of the corresponding object. Media owners can thereby track the dissemination of copies of their media objects – at least insofar as use of such objects entails communicating with the associated MGA site.

Such tracking offers a great number of opportunities, some in the area of commerce. The MGA site corresponding to the cover art of a Garth Brooks CD, for example, can provide a listing of IP addresses of persons interested in that CD. Email or promotional data objects (e.g., audio clips) can be sent to that list of addresses when a subsequent Garth Brooks CD is released.

Such tracking also opens up a new dimension of internet searching. Presently, internet search engines use a brute force approach, visiting millions of pages across the web in order to identify, for example, a dozen instances of a given photograph file. MGAs offer a shortcut to such brute force approaches. With the present technology, a search engine can find a single instance of a photograph file and, by detection of the MGA data watermarked therein, link to the corresponding MGA site. From the MGA site, the search engine can obtain a listing (if such queries are authorized) of some or all of the other sites known by the MGA site to have copies of that photograph file. (Providing such data to search engines is a commerce opportunity for such MGA sites, which may permit such access to its listing of sites only in exchange for a fee. Or the MGA site may arrange to collect a tribute payment from the search engine proprietor each time the engine responds to a user query using data collected from the MGA site.)

Many of the addresses logged by the MGA may not be publicly-accessible data stores. The search engine can check each listed address to ensure that the desired object is present and accessible before adding the address to its database.

5 Covert Tracing

Co-pending application 09/185,380 describes anti-counterfeiting technology that looks for the presence of digital data corresponding to bank note imagery in a computer system, and makes a covert record of any attempt to process such data (e.g., Scan, FileOpen, FileSave, Print, Edit, etc.). Such records are hidden from the user of the system (using, e.g., various data encryption and obscuring techniques), but authorized law enforcement officials are provided tools by which these records can be recovered. The forensic data thereby obtained may prove useful in prosecuting counterfeiters. (Knowledge that a computer may be covertly storing evidence of attempted counterfeiting actions may prove as, or more, valuable in deterring counterfeiting than the covert records themselves.)

15 The same techniques can be employed to deter unauthorized processing of audio, image, video, or content by media pirates. In one embodiment, a computer's operating system (including peripheral device drivers) monitors various data within the system (e.g., data sent to writeable storage media, or sent via a serial port or network connection, etc.) for data bearing a do-not-copy watermark. The presence of such data being sent, e.g., to a writeable disk or to a remote computer, indicates that the do-not-copy instruction has been circumvented. In such case, the operating system writes one or more covert records memorializing the activity, for possible use in criminal prosecution if the computer is lawfully seized.

25 The example just-provided is but one of many monitoring and response techniques that may be employed to deter circumvention of copy-protection or other access control systems. Generally speaking, if content data is found where it shouldn't be, or is found used as it shouldn't be used, a corresponding record should be made. (Other intervention actions can be triggered as well; covert tracing is desirably just one of several parallel responses to suspected hacking.)

Meta-Data Accessed Using Watermarks

Meta-data, in formats known as XML, SGML, and HTML, is widely used to communicate information about digital objects (e.g., author, keywords, price, rights, caption, etc.). More generally, meta-data can be thought of as any data construct which associates the name of a property (e.g., "author), with the value of the property (e.g., "Mark Twain"). Such data commonly appears in a tag format, such as the following:

```
<META NAME="author" CONTENT="Mark Twain">
```

Meta-data is commonly exchanged between server and client computers in conjunction with the digital objects to which they relate (e.g., the text of a Mark Twain book).

As detailed herein, an important application of watermarking is likewise to convey information about media – in this case embedded within the media content itself (e.g., providing unique identification, establishing some basic behaviors such as do not copy, and providing links to extended functionality).

For meta-data to be useful, it must be linked to associated content, whether in the context of a browser, application program, operating system, asset management system, search engine, etc. However, as detailed below, the content and the associated meta-tags needn't always be conveyed together.

Consider an application program or other client process that receives a watermarked media object. The watermark includes an MGA for that object (which, as noted above, may not specify an ultimate IP address). Stored at the MGA site is meta-data corresponding to the object. By linking to the MGA site identified by the object's watermark, the client computer can obtain the meta-data corresponding to the object. This data can be stored at the client computer and used just as any other meta-data, e.g., to define the local functions that should be available for use with that object (e.g., buy, search, etc.)

A particular example is an on-line catalog of stock photography. Each photograph is watermarked with MGA data. To identify the photographer, copyright date, price, telephone number, subject, etc., an application program can link to the MGA site for that photograph, and obtain the corresponding meta-data. This data can then be displayed or used as needed. Data objects of disparate formats thus can readily be handled within a single, simple application program, since the program needn't concern itself with the varying formats for the associated

meta-data (assuming the name servers provide this data in standardized format). Substantial flexibility in programming and object formatting is thereby achieved.

Returning to the internet search engine example described above, MGAs may become recognized as repositories rich in meta-data for media objects. Specialized search engines may focus their data collection around such sites, and be able to quickly identify the MGA sites corresponding to various boolean combinations of meta-tag parameters.

Asset Management/Containers

Much has been written on the topic of asset rights management. Sample patent documents include U.S. Patents 5,892,900, 5,715,403, 5,638,443, 5,634,012, 5,629,980 and laid-open European application EP 862,318. Much of the technical work is memorialized in journal articles, which can be identified by searching for relevant company names and trademarks such as IBM's Cryptolope system, Portland Software's ZipLock system, the Rights Exchange service by Softbank Net Solutions, and the DigiBox system from InterTrust Technologies.

An exemplary asset management system makes content available (e.g. from a web server, or on a new computer's hard disk) in encrypted form. Associated with the encrypted content is data identifying the content (e.g. a preview) and data specifying various rights associated with the content. If a user wants to make fuller use of the content, the user provides a charge authorization (e.g. a credit card) to the distributor, who then provides a decryption key, allowing access to the content. (Such systems are often realized using object-based technology. In such systems, the content is commonly said to be distributed in a "secure container.")

Desirably, the content should be marked (personalized/serialized) so that any illicit use of the content (after decryption) can be tracked. This marking can be performed with watermarking, which assures that the mark travels with the content wherever -- and in whatever form -- it may go. The watermarking can be effected by the distributor -- prior to dissemination of the encrypted object -- such as by encoding a UID that is associated in a database with that particular container. When access rights are granted to that container, the database record can be updated to reflect the purchaser, the purchase date, the rights granted, etc. An alternative is to include a watermark encoder in the software tool used to access (e.g. decrypt) the content. Such an encoder can embed watermark data in the content as it is released from the secure container, before it is provided to the user. The embedded data can include a UID. This UID can be

assigned by the distributor prior to disseminating the container. Alternatively, the UID can be a data string not known or created until access rights have been granted. In addition to the UID, the watermark can include other data not known to the distributor, e.g. information specific to the time(s) and manner(s) of accessing the content.

5 As noted earlier, access rights systems can be realized with watermarks without containers etc. For example, in a trusting world, copyrighted works can be freely available on the web. If a user wishes to make lawful use of the work, the user can decode its watermark to determine the work's terms and conditions of use. This may entail linking to a web site specified by the embedded watermark (directly, or through an intermediate database), which specifies the
10 desired information. The user can then arrange the necessary payment, and use the item knowing that the necessary rights have been secured.

Remote Reconfiguration of Watermark Detectors

15 In some cases, it is desirable to reconfigure watermark detectors remotely. Such functionality is desirable, for example, if a watermark system is hacked or otherwise compromised.

In accordance with this aspect of the present invention, some aspect of a watermark detector's operation is changed in response to a command. The change can take various forms. In watermark systems employing pseudo-random key data (e.g., spread spectrum spreading
20 signals), the pseudo-random signal used for detection can be changed. In systems using DFT processing, the mapping between message bits and DFT coefficients can be changed. In still other systems, the decoding can proceed as before, but the significance of one or more bits can be changed (e.g., bits that were normally interpreted as defining Field A can be interpreted as defining Field B, and vice versa). In yet other systems, the decoding can proceed as before, but
25 the response of a device to a given watermark signal can be changed. In still other systems, a set of software instructions can be re-written or re-ordered to effect a change in detector operation.

The command can be conveyed in various ways. In one embodiment, it can be a trigger bit in the watermark payload. Normally the bit has a value of "0." If the bit has a value of "1,"
30 the detector system responds by changing its operation. A trigger pattern can also be established, so that detection of a certain combination of bits in the watermark payload serves to trigger the change. Reserved states of certain data fields are examples of patterns that might be employed.

The command can also be conveyed through another channel different than the watermark channel (e.g., an SCA channel of an FM broadcast, or the sub-titling data channel of video broadcasts, or header data within an MPEG data stream, etc., etc.).

5 The change can proceed in accordance with a pre-programmed rule (e.g., codes progressing successively through a numerically or algorithmically-determined progression), or the change can proceed in accordance with data specified elsewhere in the payload of the watermark bearing the trigger bit (e.g., instead of being interpreted in normal fashion, the non-trigger bits of the detected watermark can define a new pseudo-random key data. Or the change can proceed in accordance with data conveyed in successively-presented watermark payloads, as
10 might be done in video encoding where each frame of video can convey further watermark information. (This latter arrangement is one offering a high-bandwidth re-programming channel through which, e.g., extensive firmware instructions might be transferred to the detector to replace instructions earlier stored.)

15 By such arrangements, greatly increased detector versatility and functionality can be achieved.

Conclusion

Many diverse embodiments are reviewed above – each with a unique set of features. (Still others are disclosed in the assignee's patents incorporated by reference.) This specification
20 should be construed as explicitly teaching that features illustrated in one such embodiment can generally be used in other embodiments as well. Thus, for example, a date field was not particularly discussed in connection with payload data for video watermarking. Nor were "play once" watermarks so-considered. The inclusion of a calibration signal with (or as part of) the watermark is shown in embodiments of the issued patents, but is not belabored in the above-
25 described embodiments. Likewise with "simple universal codes." The pre-stored commerce profile described in one of the foregoing embodiments is equally applicable to other embodiments as well. Likewise, the presentation of advertising was discussed in connection with one embodiment but not others, although it, too, is generally applicable. All of these concepts are familiar at Digimarc and are regarded as generally applicable throughout the work
30 expressed in Digimarc's patent disclosures. Practicality prevents an exhaustive recitation of each individual permutation and combination.

Having described and illustrated the principles of our invention with reference to illustrative embodiments, it will be apparent that the detailed arrangements can be modified in arrangement and detail without departing from such principles.

5 For example, while reference has been made to various uses of wireless, it should be understood that such reference does not just cover FM broadcast, and wireless internet networking and the like, but also includes other wireless mechanisms. Examples include cell phones and direct satellite broadcast.

10 Likewise, while certain embodiments were illustrated with a watermark payload of 100+ bits, in other systems much smaller (or sometimes larger) payloads are desirable – sometimes as small as 1-8 bits.

While the foregoing examples have each been illustrated with reference to a particular media type (e.g., video, audio, etc.), it will be recognized that the principles of each embodiment find application with the other media types as well.

15 Certain of the appliances contemplated above require user interfaces more sophisticated than are presently typical on such devices. The simplicity of the underlying audio appliance can be preserved, in many instances, by using a palmtop computer – coupled by infrared or otherwise – as a temporary user interface to the appliance. Some of the processing capability can likewise be off-loaded to an ancillary palmtop. (Palmtop is here meant to refer generally to any pocket-size programmable computing device.)

20 Unless otherwise stated, it should be understood that the digital music, video, and imagery contemplated herein is not of any particular form or format. Audio, for example, can be of various forms, both streaming and non-streaming, and of various formats (e.g. MP3, MP4, MS Audio, Windows Media Technologies, RealAudio, *.WAV, MIDI, Csound, Dolby's Advanced Audio Codec (AAC), etc.

25 To provide a comprehensive disclosure without unduly lengthening the present specification, applicants incorporate by reference the patent publications and applications cited herein.

We claim as our invention all such embodiments as may come within the scope and spirit of the following claims, and equivalents thereto.

WE CLAIM**1. A method comprising:**

5 encoding digital source material to steganographically convey plural-bit auxiliary data;
passing the encoded source material to a destination through at least one intervening
computer;
at said intervening computer, detecting encoded source material transmitted thereby; and
crediting a payment in response to said detection of the encoded source material, in
accordance with the plural-bit auxiliary data steganographically conveyed by the encoded source
10 material.

15 2. The method of claim 1 which includes decoding plural-bit auxiliary data only from source
material that has first been tested to indicate the likely presence of such auxiliary data
therein.

3. The method of claim 2 which includes testing objects by reference to an encoding attribute
that is supplemental to said encoded plural-bit auxiliary data.

20 4. The method of claim 3 in which said attribute is the presence of a characteristic signature
signal conveyed by said object.

5. The method of claim 4 in which the signature signal is a repetitive noise burst signal.

25 6. The method of claim 1 in which said transmitting includes distributing through a network of
interconnected computers.

7. The method of claim 1
reporting said detection to a location remote from detection over same network
crediting royalties based on detection

8. A method comprising:

encoding audio source material to steganographically convey plural-bit auxiliary data;
presenting the audio source material to a consumer;

decoding the audio source material as it is being presented to the consumer, to decode the
auxiliary data therefrom; and
storing data indicating the audio source material(s) presented to the consumer.

9. The method of claim 8 that includes generating a report based on the stored data, indicating
the audio source material(s) presented to the consumer.

10. The method of claim 8 which includes detecting the presented audio source material with a
microphone, and decoding the auxiliary data from a microphone output signal.

11. A method comprising:

encoding an object to steganographically convey plural-bit auxiliary data;

distributing the object beyond the control of a proprietor thereof;

thereafter, decoding the plural-bit auxiliary data from the object;

consulting a registry to determine the proprietor of the object, by reference to said decoded
plural-bit auxiliary data; and

making a payment to said proprietor.

12. The method of claim 11 that includes making said payment through the registry.

13. The method of claim 11 in which the object is a work of authorship, and the encoding adds a
generally imperceptible level of noise to the object as it is perceived by a consumer thereof.

14. The method of claim 11 in which the registry comprises a database accessible through the
internet.

15. A method of encoding a digital object, comprising:

encoding the object with a first information signal, said first information signal having relatively small information content, but permitting rapid decoding; and

5 encoding the object with a second information signal, said second information signal conveying having relatively high information content, requiring relatively more time to decode.

16. The method of claim 15 in which the first information signal is a signal indicating to decoding equipment that the object is not to be copied, and the second information signal is a signal conveying information relating to ownership of the object.

10 17. The method of claim 15 in which:

the digital object is a digital representation of music; and

the first information signal is a broadband, repetitive signal that is conveyed at a low level within said music.

15 18. The method of claim 15 in which the first and second signals are independent of each other.

19. The method of claim 15 in which the first and second signals are aspects of a combined signal.

20 20. A method of processing an object that has been steganographically encoded with first and second information signals, the first information signal having relatively small information content, the second information signal having relatively larger information content, the method comprising:

15 decoding from the object the first information signal, the relatively small information content of the first information signal permitting relatively rapid decoding;

disabling an operation of an apparatus in accordance with the decoded first information signal; and

) optionally, decoding from the object the second information signal, the relatively larger information content of the second information signal requiring relatively more time to decode, said second information signal conveying information relating to ownership of the object.

21. A method of encoding audio with a marker signal indicating a restriction on permitted copying, wherein the marker signal is characterized by being in-band, broadband, and repetitive.

5

22. A method comprising:

watermarking plural-bit binary payload data in an object;

reading the payload data from the object using a device; and

10 using the payload data read by the device in connection with a commercial transaction involving music related to said object.

23. The method of claim 11 in which the object is a poster having artwork thereon.

15 24. The method of claim 11 in which the object is a storage medium having a music video recorded thereon.

25. The method of claim 11 in which the device is a handheld, battery powered device.

20 26. A method of altering music data to steganographically insert plural bits of watermark data therein, characterized by inserting a first group of said bits for benefit of an end-user of the music data, inserting a second group of bits different than the first for benefit of an artist whose music is encoded by said music data, and inserting a third group of bits different than the first two for benefit of a distributor of the music data.

25 27. The method of claim 26 in which the first group of bits represents an internet address of a web site that may be accessed by end-users of the music data.

0 28. The method of claim 26 in which the second group of bits includes bits representing a unique identifier for the music data, permitting machine identification of the data and royalty credit to the artist.

29. The method of claim 26 in which the third group of bits represents usage restrictions to which audio appliances are responsive, thereby driving distribution of additional copies of the music data.

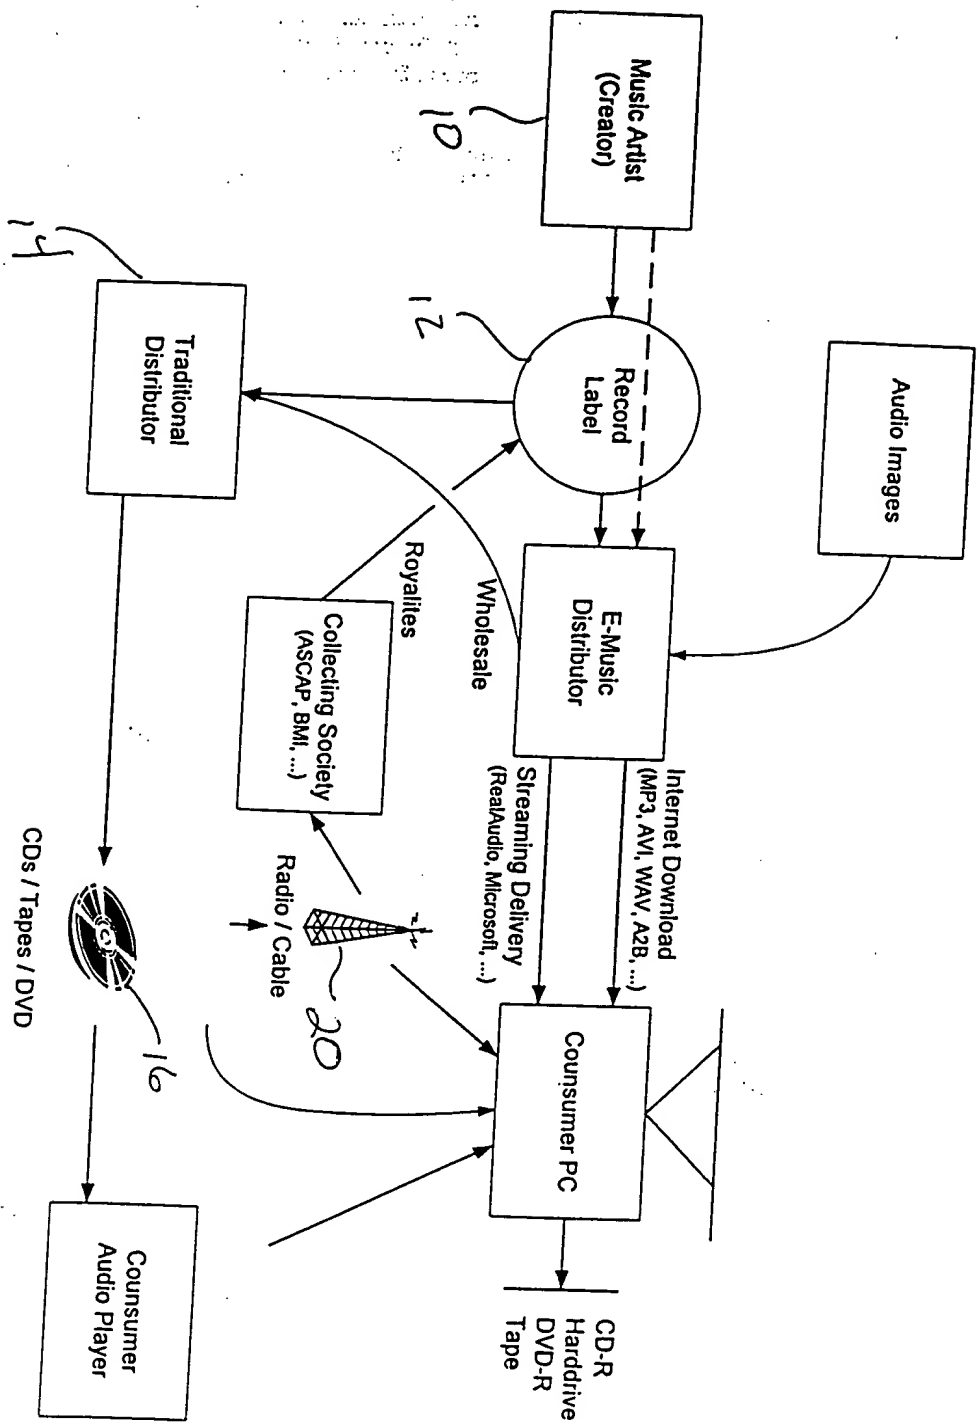


FIG. 1

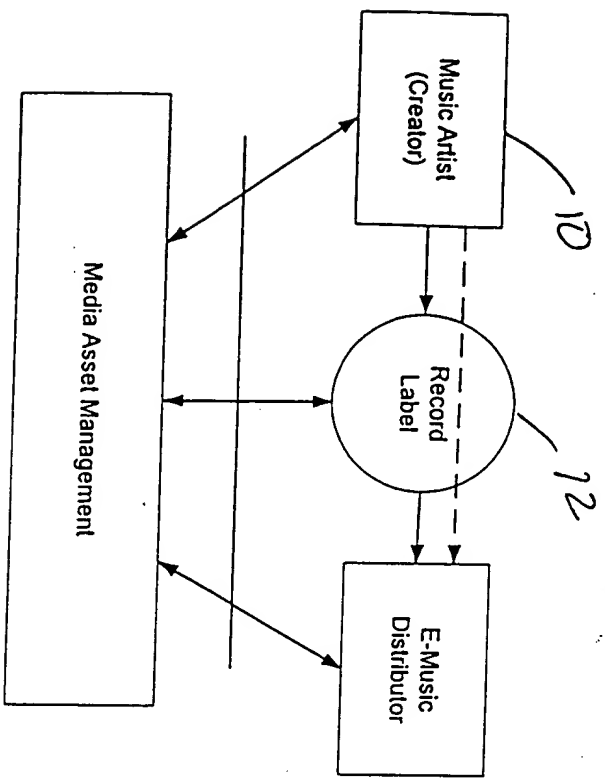


FIG. 2